Daniel Ellsberg

Notes Toward a Theory of Wishful Thinking

Under certain conditions, the hypothesis I will discuss implies behavior violating the Ramsey-Savage-Luce-Suppes postulates on decision-making under uncertainty. These conditions were met in some recent betting experiments by John Chipman; he reports violations of the sort suggested by my hypothesis. Since this constitutes the only experimental test of my hypothesis in the field of betting behavior (it is supported by considerable experimental evidence in the field of verbal expectations, forecasting and guessing behavior), I will discuss his results and my interpretation of them. (Chipman's interpretation is different, incidentally; further experiments will be needed to test which generalization is more useful). I will then spell out my hypothesis in full and draw attention to some remaining empirical implications that conflict with current theories.

I

Let (aAb) denote an act which will have consequence _a_ if event A occurs, _b_ if A does not occur; we will call such an act a _gamble_. The Luce-Savage-Ramsey-Suppes axioms are restrictions on a subject's preferences (choices) among gambles of this typical form. Among the restrictions implied are the following:
(a, b are outcomes such that a is preferred to b; X, Y,Z are events)

a) If he prefers aXb to aYb, then he prefers bYa to bXa.

b) If he is indifferent between aXb and bXa, and between aYb and bYa, then he is indifferent between aXb and aYb.

c) If he prefers aX_1b to aY_1b and aX_2b to aY_2b, he prefers $a(X_1$ or $\cancel{X_1})b$ to $a(\overset{Y}{\underset{x}{X}}_2$ or $Y_2)b$.

X_1, X_2 disjoint

The basic inference in the ~~current~~ Ramsey approach is that if the subject prefers a to b, and prefers aXb to aYb, then he considers event X "more likely than" event Y. (It is also assumed that the gamble aA'b is equivalent--and indifferent--to the gamble bAa, where ~~ai~~ A' denotes "not-A"). With this interpretation, the above conditions correspond to the following:

a) If X is more likely than Y, then not-Y (Y') is more likely than not-X (X').

b) If X and X' are equally likely, and Y and Y' are equally likely, then X and Y are equally likely (likewise X and Y', X' and Y).

c) If X_1 is more likely than Y_1, and X_2 is more likely than Y_2, then the union of X_1 and Y_1 is more likely than the union of X_2 and Y_2.

Clearly, these are characteristics that the relationship "more likely than" (as defined operationally above) must have if it is to be interpreted as a ~~qualitiat~~ qualitative probability. (Stronger restrictions on choices are required--and supplied by the authors cited--for us to infer a numerical probability from choices among gambles. However, I will examine only those conditions that underly a qualitative probability relationship among events; as seen above, this implies a relationship with more structure than a weak ordering of events).

I propose the following experiment. A subject confronts two urns containing red and black balls, from one of which he will draw a ball at random. In the case of urn I, he knows nothing about the ratio of red to black balls; it may be anything from 0 to 1. Let R_1

represent the event that, drawing from urn I, he draws a red ball ($R_1' \equiv B_1$ is the complementary event of a black ball drawn). Let a and b denote payoffs of \$2 and \$0, respectively. If he is asked to choose between aR_1b and bR_1a (i.e., between aR_1b and $aR_1'b$), which will he prefer? Let us suppose that he is indifferent between the two gambles. Ramsey, et. al., would infer that the subjective probabilitëes for R_1 and R_1' are equal; we assign the number $\frac{1}{2}$ to each.

The subject then counts the balls in urn II and discovers that there are 50 red and 50 black balls, in random order. If R_2 is the event that,axx drawing from urn II, he picks a red ball, let us suppose that again he is indifferent between betting on red and betting on black. I.e., he is indifferent between the gambles aR_2b and bR_2a. Again we assign probabilities of $\frac{1}{2}$ to R_2 and R_2'.

Now we offer him a choice between aR_1b and aR_2b. He can bet on red, drawing from urn I, or (the same amount) on red, drawing from urn II. Will he now be indifferent? My intuition suggests strongly that he may not. If not, he is violating condition (b) above.

This is the point on which there is some experimental evidence. In Chipman's experiments, the subject knew the population frequency in "urn II"; he observed a small, allegedly random sample from "urn I." In one case, the sample from urn I consisted of five red and five black balls ; the population inurn II was fifty red and fifty black. (I am adapting his experimental conditions to my hypothetical ones. Chipman actually used mathhboxes, and the heads and stems of matches; smaller payoffs; and a stochastic definition of preference and indifference. I will discuss the latter below).

In this case, hës subjects were indifferent between aR_1b and

bR_1a, and also between aR_2b and bR_2a. But to a significant degree, they were not indifferent between betting on red from the sampled urn and red from the known-population urn (i.e., between aR_1b and aR_2b). As it happened, they strongly preferred urn II; they chose to draw from the urn with the known 50-50 distribution. (I suspect that the preference might run the other way under conditions; possibly with larger amounts of money. The point I wish to note is the lack of indifference). They violated condition (b).

Again, suppose that we observe ~~began by observing~~ that a subject pre-ferred to bet on red from urn II rather than red from urn I; he preferred aR_2b to aR_1b (as Chipman's subjects did). In Ramsey's interpretation, this would mean that he regarded R_2 as "more likely than" R_1. By condition (a), then, he must regard B_1 as more likely than B_2. At least, he must act "as though" he did; he must prefer aB_1b to aB_2b. But will he? It seems likely to me that a subject who prefers to bet on R_2 rather than R_1 will also prefer to bet on B_2 rather than B_1 (by considerations of symmetry; ~~remember that he~~ I assume that he will act as though R_1 and B_1, and R_2 and B_2, are equally likely). If his choices reveal only his estimates of likelihood, then he regards R_2 as more likely than R_1, yet also regards not-R_2 more likely than not-R_1! Such relationships between events do not have the structure of a qualitative probability. (Specifically, they violate condition (a), above).

Certain other results of Chipman's experiment have a bearing on condition (c). Suppose that urn I contains red, yellow and black balls, and a random sample consists ~~to~~ of 3R-3Y-3B. Suppose that

urn II contains exactly 30B, 30Y and 30R. First the subject chooses between aR_1b and aR_2b; then between aY_1b and aY_2b. Chipman found that when the frequency ratios in the sampled and the "known" urn were equal and the proportion of "favorable" outcomes was less than $\frac{1}{2}$, subjects preferred to draw from urn I. In this case, this means they would act as though R_1 were ~~less~~ more likely than R_2 and Y_1 more likely than Y_2.

Finally, the subject chooses between $a(R_1$ or $Y_1)b$ and $a(R_2$ or $Y_2)b$. In other words, he "wins" if either a red ~~and~~ or a yellow ball comes up. In Chipman's experiment, when the proportion of "favorable" outcomes was greater than $\frac{1}{2}$, subjects chose to draw from urn II. Thus, they would choose $a(R_2$ or $Y_2)b$, indicating that the union of R_2 and Y_2 was more likely than the union of R_1 and Y_1: in violation of condition (c).

In Luce's approach, discrimination between likelihoods and between payoffs is assumed to be imperfect, and the Savage-Ramsey axioms are replaced by restrictions on the statistical distribution of choices. The restriction corresponding, for example, to condition (a) above is his Lemma 5, which states that the statistical frequency with which aAb is preferred to aBb must equal the frequency with which aB'b is preferred to aA'b. In other words, B' must be judged more likely than A' with the same frequency that A is ~~preferred~~ judged more likely than B. This condition is violated by the hypothetical results suggested above, and by Chipman's actual results for the 50-50 case.

III

The Savage-Ramsey-Luce axioms give an operational meaning to the statement that individuals act "as though" they assigned qualitative or numerical probabilities to events. In this sense, Chipman's subjects do not act as though they assigned such probabilities: nor do the hypothetical subjects described above. What <u>are</u> they doing? Is there a pattern in their behavior that can be axiomatised by a different set of postulates?

One approach to an answer would be to search for violations of specific axioms, in hopes that the separate axioms will provide useful categories for distinct types of violations. It may turn out that an important class of behavior can be "explained" by abandoning or modifying one particular axiom.

Instead, I will begin with an alternative set of hypotheses to explain the behavior. They happen to be derived, not from consideration of actions violating the Savage axioms, but from positive results of psychological experiments in fields related to (but not including) betting: guessing, forecasting, verbal expectations. To present these results here would, I believe, make the following hypotheses far more "plausible." However, I will restrict myself to outlining ~~the~~ their empirical implications, drawing attention to the fact that they do imply violations of the Savage-Luce axioms and that they are consistent with Chipman's and the hypothetical results.

In the first sentence of their book, "Decision Making," Suppes and Davidson emphasize the two variables that dominate the Ramsey-Savage-Luce approach:

"When an individual chooses between alternatives which involve uncertain outcomes it seems clear that at least two factors enter the decision: the degree to which the possible outcomes are desired relative to one another, and the degree to whichthe outcomes are deemed probable."

The caution, "at least," goes unnoticed in the remainder of the analysis, which allows roles only for the two facters, relative desirability of payoffs and relative likelihood of events. I propose to introduce two additional variables: ambiguity of events, and importance of the payoffs. I will postpone a detailed operational definition of these concepts, but their meaning can be made clear in specific contexts. "Ambiguity" is associated with the quantity, type and "unanimity" of information about the likelihood of the events. For example, if the ~~distribu~~ information about all the events in a set of gambles is in the form of sample-distributions, then "ambiguity" may be closely (inversely) related to the size of the sample. Indeed, Chipman (like Georgescu-Roegen) introduces sample-size as a variable for virtually the same purposes that I use "ambiguity," though the hypotheses he bases on it differ from mine. Since I will be describing Chipman-type experiments, sample-size will serve as a working definition of ambiguity. (I don't favor Chipman's concept in general because: a) information about many events can't conveniently be described as a sample-distribution; b) emphasis on sample-size suggests emphasis on "quantity" of information. It may be useful to regard "ambiguity" as high even though there is an ample "quantity" of information, when there are questions of reliabilit and particularly when there is conflicting opinion and evidence.) "Importance" is related to the level and sign of the payoffs.

I find that both Luce and Savage mention similar concepts, ~~though they give them no role in their approach~~. Luce discusses the

"fuzziness" of probabilities associated with everyday events. He assumes that resulting inconsistencies of choice will still conform to his stochastic axioms. Savage refers to the "aura of vagueness" that surrounds many judgments of probability; unlike Luce, he suggests that some such factor might conceivably underlie such behavior as "minimaxing regret," which does violate the axioms, but he indicates that he does not regard such behavior as empirically significant. I will invoke a notion similar to "fuzziness" or "vagueness" specifically to explain significant violations of the axioms.

Let us imagine an experiment in which the variable "ambiguity" is experimentally controlled. Assume that there is a critical level of ambiguity so that we can classify events either as ambiguous or non-ambiguous. Let us suppose that the relative likelihood of certain <u>classes</u> of events are unambiguous with respect to each other, but that the relative likelihoods of events within a given class may be ambiguous (though they need not be equal; ambiguity is not the same notion as likelihood).

For example, statistical analysis of the results of horse races reveals that the "favorite" wins between 1/3 and 1/2 the time; it does not reveal much else. We may imagine that a particular bettor, knowing this, assigns an "unambiguous" likelihood of 1/3 to the favorite (knowing nothing about the horses except the amounts bet) and 2/3 to the pack of non-favorites: but that the relative likelihoods of the individual non-favorites are ambiguous (not necessarily equal).

Chipman's experiment can be modified slightly to approach this situation. Suppose urn I contains ~~an unknown~~ 100 red and black balls in unknown proportion (or sampled: 5R-5B); urn II contains 100 yellow and blue balls, fifty of each. The two urns are dumped into one large

urn and contents randomized. We may suppose that he assigns equal probabilities to (Red or Black) and (Yellow or Blue), and that he regards Red as equally likely to Black, Yellow to Blue. That is, he is equally willing to bet on either member of these pairs of events. We can represent such choices by the following matrix:

		? 100	?	50 100	50
		Red	Black	Yellow	Blue
	I	a	a	b	b
	II	b	b	a	a
a > b	III	a	b	b	b
	IV	b	a	b	b
	V	b	b	a	b
	VI	b	b	b	a

We assume, then, that he happens to be indifferent between actions I and II, between III and IV, and between V and VI. Now, is he indifferent between, say, IV and V? Is he indifferent between betting on Black or Yellow? He knows that there are exactly 50 yellow balls, out of 200; of the number of black balls, he knows (only) that there are between 0 and 50.

Savage's Postulates 1 and 2, the simple ordering of acts and the "Sure-thing principle," together imply that actions III, IV, V and VI are all mutually indifferent, on the above premises (likewise that abab and baba, aabb and abba, abaa and aaba, etc., are ~~equally likely~~ indifferent). The intuitive plausibility of this result is not overwhelming. Chipman's results alone weigh heavily against it. Yet from the usual point of view, to deny indifference involves paradox. How can it be, for example, that a sane subject prefers V to IV, VI to III, yet is indifferent between I and II? Under the usual interpretation, this would indicate that he regarded Yellow as more likely than Black, Blue more likely than Red, yet (Yellow or Blue) equally likely to (Red or Black)! Yet I maintain that I would

actually expect some such result.

The paradox persists only so long as we try to explain the behavior on the basis of two factors alone, relative likelihood and relative preference.(the latter held constant here). The notion of ambiguity as a separate variable makes immediate sense here. Intuitively, the likelihoods of the events (Red or Black), (Yellow or Blue), Yellow, Blue are unambiguous ~~(they happen here to be all equal)~~; the likelihoods of Red, Black, or such events as (Red or Yellow), (Black or Blue) are ambiguous, though they have definite boundaries. By my hypothesis, a choice between betting on Black or on Yellow "reveals" ~~not merely~~ (is based upon) not merely a measure of relative likelihood of the two events but ~~upon~~ a measure of differing ambiguity. (At this point, the reader may notice a similarity of role between this concept and the old one of "confidence." The roles are indeed similar, but I have implicitly defined "ambiguity" so far in objective terms, as a property of the information available about an event. A thorough discussion of the operational definition of ambiguity would take up too much space here, but ~~it~~ the term need not be defined subjectively, or even behavioristically).

what is that

<u>How</u> will this new dimension influence behavior? I believe that this will vary between individuals and between different situations. However, there is considerable evidence in favor of one or two particular hypotheses. For one, a great deal of advice about "conservative" policy implies a dependence of expectations on payoffs, in the direction of "expecting the worst." As Fellner puts it, "There is a strong presumption that the possibility of less favorable outcomes is weighted heavier than the contrary." Whether or not such advice is "rational," it is reasonable to expect that it is sometimes

followed . Let us suppose that the tendency appears primarily with respect to ambiguous events, and that it takes an extreme form. Assume that the full likelihood of a subclass of ambiguous events is attached, given any act, to the worst possible outcome under that set of events. That is, with respect to this set of events, the subject "minimaxes." (This extreme case is meant to serve only as an approximation to actual "conservative" or "pessimistic" behavior). But with respect to unambiguous events (and the subclass of ambiguous events may, as a class, be an unambiguous event) he does not minimax; he may indeed obey the Savage axioms.

Such a man will violate Savage's Postulate 2, the Sure-thing Principle. Consider the horse-race example, where the relative likelihood of Favorite versus All Non-favorites is unambiguous (1/3-2/3) but the sub events within the subclass All-non-favorite are ambiguous. Represent this in terms of our previous examples as an urn containing 30 Red balls and 60 balls in an unknown proportion of Black to Yellow.

	Red	Black	Yellow
I	a	b	b
II	b	a	b
III	a	b	a
IV	b	a	a

By Postulate 2, if I is preferred to II, then III must be preferred to IV (since the two pairs differ only by the value of a constant column). But for our conservative, the gambles reduce to the following:

	Red	(Black or Yellow)
I	a	b
II	b	b
III	a	b
IV	b	a

He will prefer I to II, but since (Black or Yellow) is unambiguously more likely than Red, he will prefer IV to III.

Since Savage believes strongly that behavior violating the Sure-thing Principle is "unreasonable," in the sense that a reasonable man is unlikely to persist in such behavior if he thinks about it for a while, he might ask such a subject if he could explain his choices. He might get some such reply as this:

"I choose according to the mathematical expectationsof payoff, where I know them. The mathematical expectation of action I is 1/3a + 2/3b; the mathematical expectation of II may lie anywhere between 2/3 + 1/3b and b. I like to be on the safe side; in fact, when I don't have any reason to expect one thing or another, I expect the worst. So I assume the mathematical expectation of II is b, and I prefer I.

But the mathematical expectation of IV is 1/3b + 2/3a, so I prefer it to III, whose mathematical expectation may lie anywhere between 1/3a + 2/3b to a."

not convincing

I don't mean to imply that a typical subject would be so self-conscious or sophistocated. But I think that a man who argued this way would not necessarily be unreasonable, and he might not change his ways even after Savage had pointed out his violation of Postulate 2. This is to cast doubt on the universality of Postulate 2 even as a normative principle. However, I bring this up simply to lend "plausibility" to the hypothesized behavior as an empirical pattern.

Actually, the bulk of the psychological evidence points in the direction of the opposite pattern: towards "wishful thinking," maximaxing instead of minimaxing. Results of experiments on guessing and forecasting indicate that subjects usually tend to emphasize the hopeful, favorable possibilities in an ambiguous situation. But this tendency is not equally strong under all circumstances; it is influenced strongly by the degree of ambiguity, and by the other variable I have bypassed so far, the importance of the decision. (We will assume that all decisions considered are important enough to merit careful evaluation; otherwise, their importance may vary). Under conditions

to here and

of high ambiguity and high importance, wishful thinking seems to be a very significant empirical phenomenon even among intelligent people.

One psychologist, McGregor, has suggested an explanation of the fact that it is often difficult, even by a conscious effort, to evade tendencies to wishful thinking. If the ambiguity of a situation consists in the abundance of conflicting forecasts and indications, the decision-maker is forced to select from the mass of conflicting evidence scraps of information on which to base his expectations and decisions. Experiments suggest that it is very hard to resist the tendency to select the hopeful indications, the opinions and evidence that favor the more desirable events (i.e., those associated with favorable outcomes). This will be more marked the more "important" the decision: the more strongly one desires or fears a particular outcome. Can we perhaps see this process, even at the very highest level of decision-making, in the evaluation of intelligence reports?

Thus, there may be a great quantity of information yet the situation be inescapably ambiguous (here is an objection to the notion of "sample-size" as an alternative to "ambiguity"). If the situation is important enough, a relatively small trace of ambiguity (rarely absent) may suffice as a basis for wishful thinking. It was possible for the Administration to disbelieve the reports of the CIA for the True? last two years. On the other hand, some events are pretty unambiguous. No one disbelieves in the existence of Sputnik, much as ~~they~~ some would like to.

Wishful thinking, of course, gives rise to the same types of violation as "pessimism." In the horse-race example above, he will

prefer II to I, yet prefer III to IV, violating Postulate 2. He is less likely to rationalize his behavior explicitly, but he might be willing to admit that where ambiguity gave him a reasonable excuse to "hope for the best," (actions III and II being ambiguous), he took it.

The same individual might appear conservative under some circumstances and wishful under others. The entire pattern of behavior of Chipman's subjects could be described by the following hypothesis: When the likelihoods of winning are about the same in the sampled urn and known urn (i.e., when the proportion of favorable outcomes in the sample and in the known population are about the same) and equal to 50% or better, the bettor is "conservative"; he chooses the known population. But when the indicated chance of winning is less than 50%, he picks the sampled urn; he chooses ambiguity, like the optimist, the wishful thinker.

One might say, when "knowledge" is pleasant, ~~he~~ this subject chooses knowledge; when "knowledge" is unpleasant, he chooses "ambiguity." When the known population offers a better than 50-50 chance of winning, there is no pressure, no "need" for wishful thinking; nor temptation to overreach for ~~a~~ still better odds. But when the known population offers, at best, the likelihood of loss, the ambiguous gamble looks good. Although he doesn't know it to be better, he doesn't know that it is not; it offers him a reasonable basis for hope. Of course, he is violating Postulates 1 and 2 in both phases.

Such an argument is heuristic, but I believe the hypotheses it suggests for testing will be fruitful. Chipman himself offers several different interpretations of his results, but none of them is able to account for the choices in the 50-50 cases. Hence, he has refrained from generalizing on the basis of those results, as I have done.

Introducing the variable "ambiguity" by itself seems to account for violations of Savage's Postulates 1 and 2. By the Sure-thing Principle (Post. 2), changes in a constant column should not affect preference among actions; but, as we have seen, it may change the ambiguity of outcome of one action (if all the events in a subclass of ambiguous events have the <u>same</u> outcome, the outcome of an action with respect to that subclass of events is no longer ambiguous), hen ce its place in the preference ordering. However, in both the psychological and economic literature referred to, it is stressed that the effect of ambiguity (whether it be wishful thinking or conservatism) depends ~~strongly~~ on the <u>importance</u> of the decision, being stronger the more important the decision. The effect of this new dimension of choice is to introduce violations of Postulate 4, the axiom specifying the independence of probabilities and payoffs.

In the Ramsey-Savage approach, the basic operation is to ~~infm~~ observe on which of two events the subject would prefer to stake a given prize; they infer that he considers this event more likely. Postulate 4 requires that this preference be independent of the particular prize offered. Thus, for any a,b,c,d such that $a > b$, $c > d$, if aXb is preferred (indifferent) to aYb, then cXd must be preferred (indifferent) to cYd.

Suppose that urn I containing 100 red or black balls in unknown proportion yields the sample, 5 red, 4 black; suppose that urn II contains exactly 49 red, 51 black. Though the event R_1 (red ball out of urn I) is more ambiguous than R_2, ~~I~~ a particular subject may prefer to bet a dollar on R_1 rather than R_2. That is, if a,b denote payoffs of \$1 and \$0, ~~I~~ he may prefer aR_1b to aR_2b. But if c,d are payoffs of \$40 and -\$30, I hypothesize that the same subject may prefer ~~aR~~ cR_2d to cR_1d. This would violate the independence postulate, precluding the unqualified inference that R_1 is subjectively more likely than R_2 (or vice versa).

Such preferences need not depend on any conscious rationale, yet again this particular subject might be willing to comment on his choices: "There is a basis for believing that the proportion of red balls is higher in urn I; anyway, I'm prepared to back that possibility with a dollar or two. But when I might lose \$30...after all, I don't know what the proportion __is__ in urn I; it might be better than 50-50, it might be a lot worse, there isn't much to go on; the safe thing seems to be to assume that it's worse than it looks. I'd rather not bet at all, but if I have to, at least I know that urn II gives me ~~better than~~ about a 50-50 chance."

Supposing that many fairly reasonable people do think and act this way (whether or not they should), how can we best describe and predict their behavior? Not within the Savage-Ramsey framework, for it violates Postulate 4. (It also violates Chipman's "lexicographic" hypothesis, in which differences in sample-size ambiguity are influential only when sample distributions are equivalent). Likewise, it would violate Luce's corresponding Axiom 2, the "decomposition axiom" which postulates that likelihood discrimination is statistically independent

of preference discrimination. If outcome a is always preferred to outcome b, then the distribution of choices between the gambles aXb and aYb will reflect only the imperfect discrimination between the likelihoods of X and Y; by Axiom 2, this distribution should be independent of the actual values of a and b. If P(aXb, aYb) is the probability, directly estimated from observation of choices, that aXb will be chosen over aYb, then $P(aXb, aYb) = P(cXd, cYd)$, for any c,d, such that $c > d$.

For the subject above, we can assume that $P(aR_1b, aR_2b) > \frac{1}{2}$, and $P(cR_1d, cR_1d) < \frac{1}{2}$. Incidentally, with different payoffs and a different ("wishful") subject, we might expect a similar phenomenon but with the preferences in the "important" decision reversed.

The following hypothesis, however, may serve to explain such behavior. The subject may assign more likelihood to ambiguous events with relatively bad outcomes when those outcomes are "very bad" than he assigns to the same events when the level of outcomes is less significant. What is the empirical meaning of this? Precisely that he will violate the specified axioms in the manner described. The man described acts "as though" B_1 (the "unfavorable" event) were less likely than B_2 when the stakes are low (though still not trivial); but when the stakes are high, he acts "as though" B_1 were more likely than B_2. In other words, when the stakes are high he acts as though his expectations were influenced by the payoffs; he violates Savage's Postulate 4, Luce's Axiom 2. If the influence were "conservative," he might consciously follow this policy and defend it; "wishful thinking," on the other hand, might be unconscious, even inadvertent, yet empirically is possibly even more common.

A last example may illustrate the type of phenomenon I am considering. In our original example, the subject was presumed to be indifferent between aR_1b and bR_1a, and between aR_2b and bR_2a, where R_1 is a red ball from an urn with unknown population and R_2 is a red ball from an urn with a known 50-50 distribution. (In a different context, R_2 might correspond to "Heads" on the next flip of a fair coin, and R_1 to "Rain in London on April 10, 1960"). Now, suppose the subject is indifferent between the gamble aR_2b and the certain outcome, c. Will he be indifferent between aR_1b and the same outcome, c?

The Savage postulates imply that he must be indifferent. The probabilities are the same, the payoffs are the same; there is no apparent basis for supposing that the "certainty equivalent" of aR_1b is any different from the certainty equivalent of aR_2b. Yet it should be clear by now that I would expect the two certainty-equivalents to be different, especially if the stakes (a,b) are high. To explain the difference, I would point to the difference in ambiguity (as I have used the term) of the two gambles. For the conservative man (who would prefer to bet on R_2 than on R_1), the less ambiguous event carries a "premium"; its certainty equivalent is higher than that of the more ambiguous event. But the "wishful thinker," the "optimist," puts the premium on the gamble that gives him a permit to hope.